GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union
NextGenerationEU

German Research Center for Artificial Intelligence

AnoMed-Seminar

Towards Privacy and Utility in Tourette Tic Detection Through Pretraining Based on Publicly Available Video Data of Healthy Subjects

Nele Sophie Brügge, Esfandiar Mohammadi, Alexander Münchau,
Tobias Bäumer, Christian Frings, Christian Beste, Veit Roessner, Heinz Handels

UNIVERSITÄT TRIER
TECHNISCHE UNIVERSITÄT DRESDEN
UNIVERSITÄT ZU LÜBECK
SHANDONG NORMAL UNIVERSITY

Medical Video Analysis in AnoMed

GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung
Finanziert von der Europäischen Union
NextGenerationEU
German Research Center for Artificial Intelligence

- Example dataset: **Toronto NeuroFace [1]**
  - Publicly available dataset
  - **Videos of oro-facial gestures performed by subjects with oro-facial impairment due to neurological disorders** with amyotrophic lateral sclerosis (ALS) and stroke.
  - Annotations: Clinical scores, facial landmarks

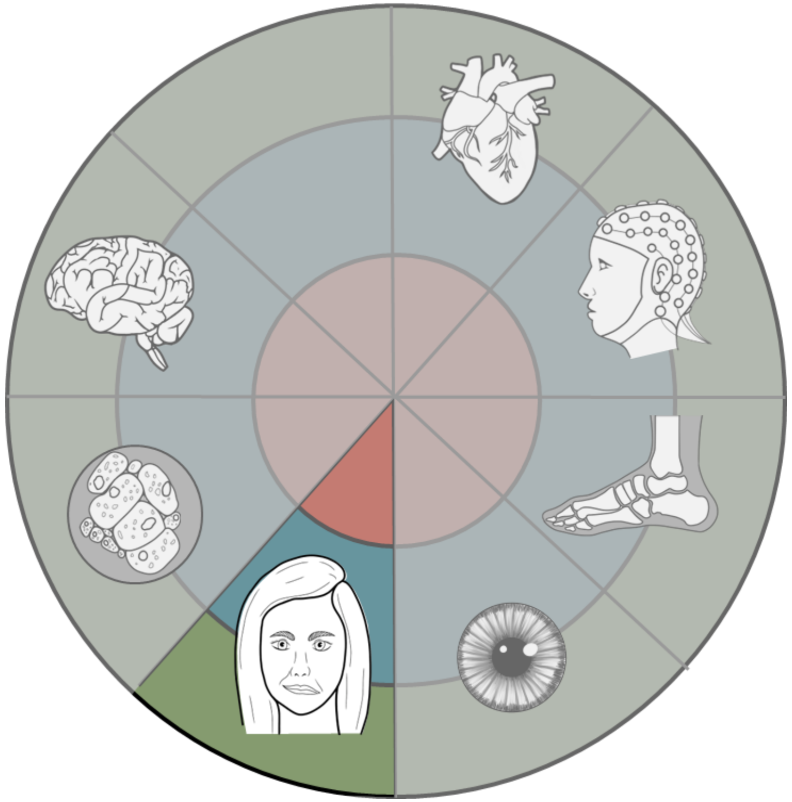

[1] Bandini, Andrea et al. "A New Dataset for Facial Motion Analysis in Individuals With Neurological Disorders." *IEEE journal of biomedical and health informatics* vol. 25,4 (2021): 1111-1119. doi:10.1109/JBHI.2020.3019242

Medical Video Analysis in AnoMed

GEFÖRDERT VOM Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union NextGenerationEU

German Research Center for Artificial Intelligence

- Example dataset: **Toronto NeuroFace [1]**
  - Publicly available dataset
  - **Videos of oro-facial gestures performed by subjects with oro-facial impairment due to neurological disorders** with amyotrophic lateral sclerosis (ALS) and stroke.
  - Annotations: Clinical scores, facial landmarks

➡ Why should medical videos be analyzed with machine learning?

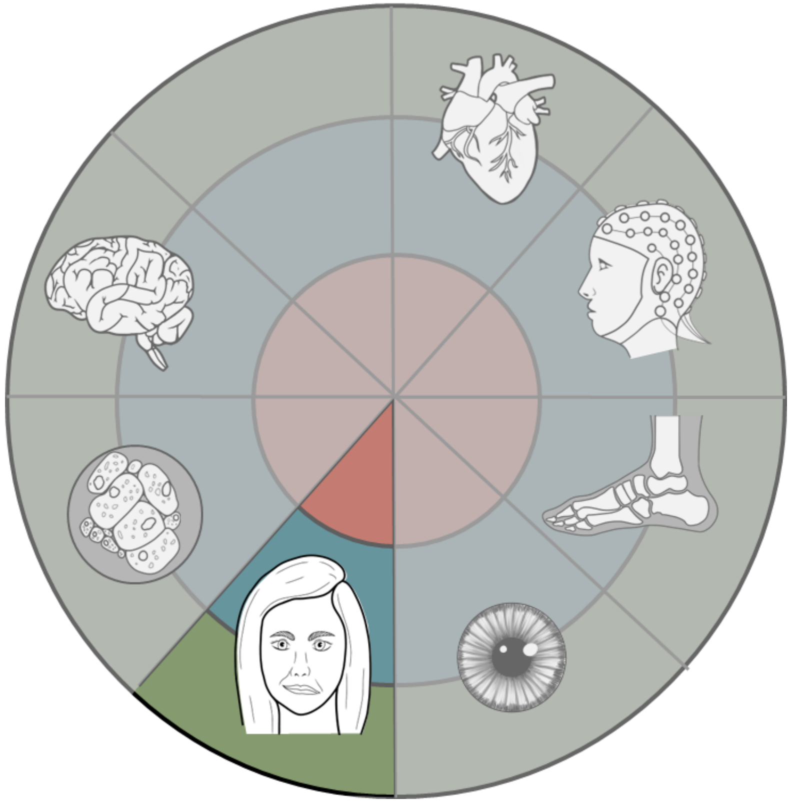

[1] Bandini, Andrea et al. "A New Dataset for Facial Motion Analysis in Individuals With Neurological Disorders." *IEEE journal of biomedical and health informatics* vol. 25,4 (2021): 1111-1119. doi:10.1109/JBHI.2020.3019242

Medical Video Analysis

GEFÖRDERT VOM Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union NextGenerationEU

German Research Center for Artificial Intelligence

Examples:

- **Gait analysis**, e.g. when using prostheses, after injuries or accidents
- **Neurological and motor disorders** (Developmental coordination disorders, stereotypic movement disorders, tic disorders)
- **Mental diseases**

[1]

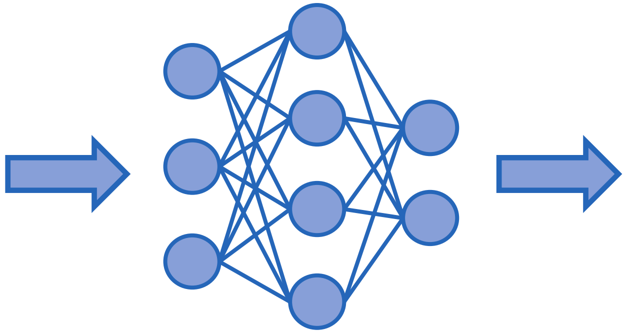

- Therapeutic monitoring
- Diagnostic support
- Treatment planning
- Disease understanding

[1] Photo: from P. Barros, N. Churamani, E. Lakomkin, H. Siqueira, A. Sutherland, and S. Wermter, "The OMG-Emotion Behavior Dataset," in 2018 International Joint Conference on Neural Networks (IJCNN), Pages: 7, Rio de Janeiro, Brazil: IEEE, Jul. 2018, pp. 1408–1414.

Motivation – Tic Detection

GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union
NextGenerationEU

German Research Center for Artificial Intelligence

Gilles de la Tourette Syndrome

- Neurodevelopmental disorder
- Prevalence: 0.3 – 0.9 %

Prior work

- Limited applicability: Additional sensors, subject-specific, healthy subjects
- No consideration of data privacy

[1]

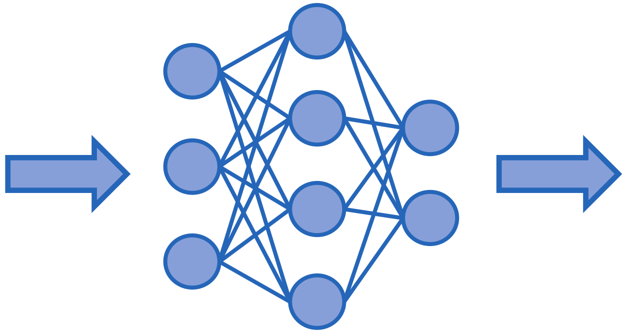

Therapeutic monitoring

Diagnostic support

Treatment planning

Disease understanding

Assessment of tic severity and frequency

[1] Photo: from P. Barros, N. Churamani, E. Lakomkin, H. Siqueira, A. Sutherland, and S. Wermter, “The OMG-Emotion Behavior Dataset,” in 2018 International Joint Conference on Neural Networks (IJCNN), Pages: 7, Rio de Janeiro, Brazil: IEEE, Jul. 2018, pp. 1408–1414.

Privacy-Preserving Machine-Learning

GEFÖRDERT VOM Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union NextGenerationEU

German Research Center for Artificial Intelligence

Ideally: Encode **general patterns** rather than facts about specific training examples
By default: Machine learning models do not learn to ignore these specifics

Attack examples:

- Reconstruction attack [3]
- Membership inference attack [4]

Challenges in Medical Video Analysis:

- High-dimensional input data
- Small datasets
- Rare diseases and special cases

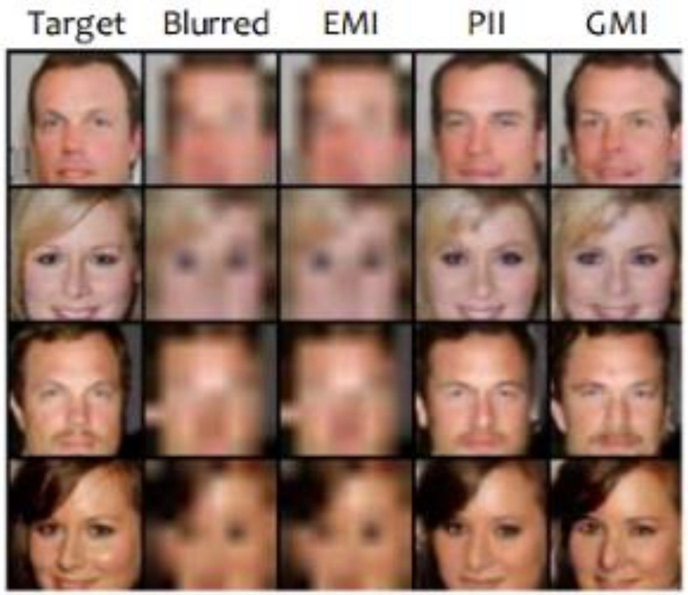

[2]

[3] Zhang, Yuheng et al. "The Secret Revealer: Generative Model-Inversion Attacks Against Deep Neural Networks." IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 250-258.
[4] R. Shokri, M. Stronati, C. Song and V. Shmatikov, "Membership Inference Attacks Against Machine Learning Models," 2017 IEEE Symposium on Security and Privacy (SP), San Jose, CA, USA, 2017, pp. 3-18.

Privacy-Preserving Machine-Learning

GEFÖRDERT VOM Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union NextGenerationEU

German Research Center for Artificial Intelligence

Ideally: Encode **general patterns** rather than facts about specific training examples
By default: Machine learning models do not learn to ignore these specifics

Attack examples:

- Reconstruction attack [3]
- Membership inference attack [4]

Challenges in Medical Video Analysis:

- High-dimensional input data
- Small datasets
- Rare diseases and special cases

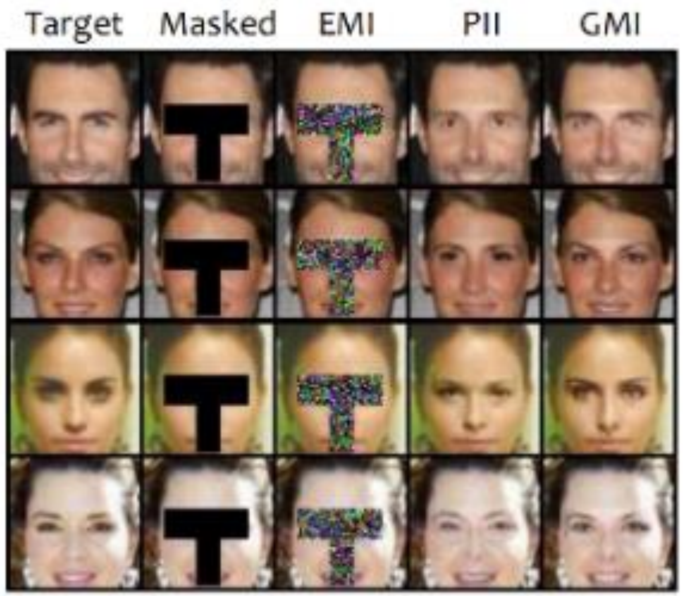

[2]

[3] Zhang, Yuheng et al. "The Secret Revealer: Generative Model-Inversion Attacks Against Deep Neural Networks." IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 250-258.
[4] R. Shokri, M. Stronati, C. Song and V. Shmatikov, "Membership Inference Attacks Against Machine Learning Models," 2017 IEEE Symposium on Security and Privacy (SP), San Jose, CA, USA, 2017, pp. 3-18.

Amplifying Membership Inference Attacks

GEFÖRDERT VOM Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union NextGenerationEU

German Research Center for Artificial Intelligence

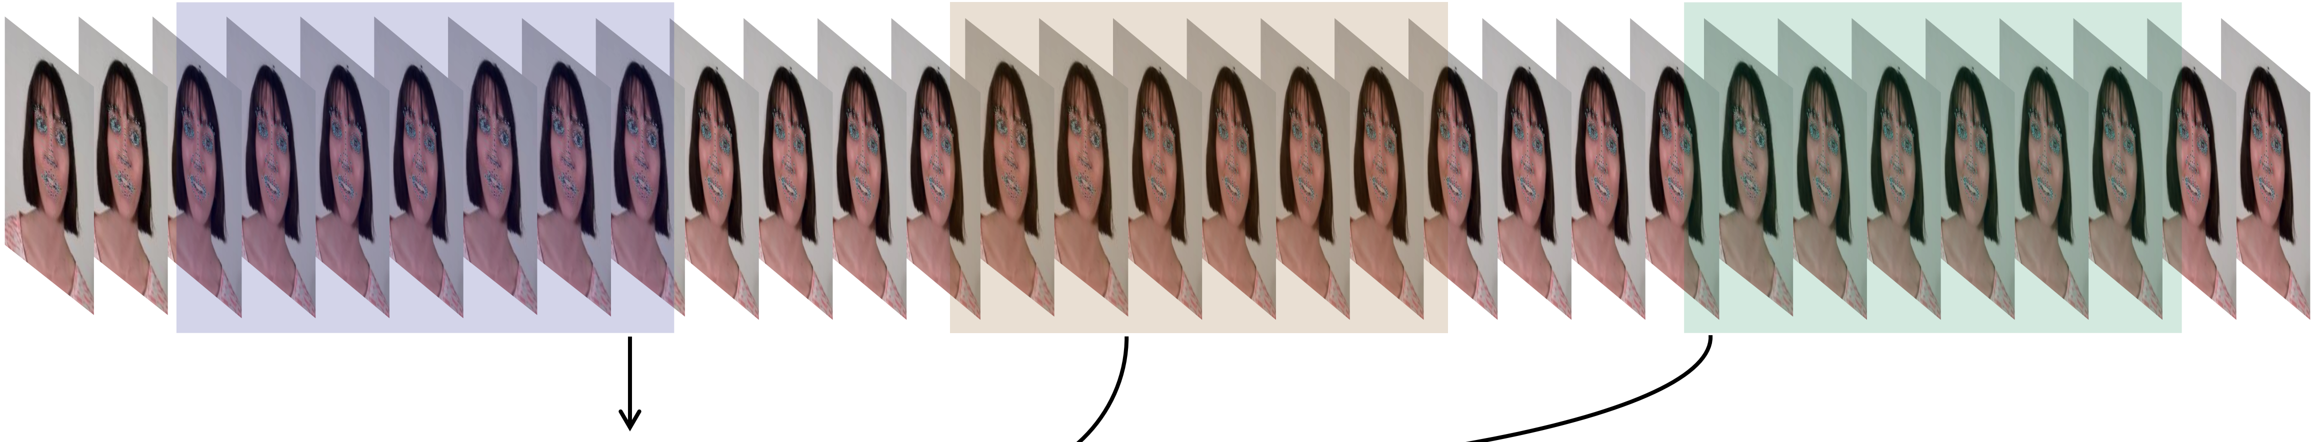

Amplification of snippet-based attacks
using Chernoff-Hoeffding bound

$$\Pr\left[\frac{1}{n}\sum X_i > p + \delta\right] \leq \exp\left(-\frac{\delta^2 n}{2p(1-p)}\right)$$ [1, 2]

n: Number of random variables
X_i: Random variables in {0,1}
p: Expectation value of X_i
δ: Deviation from expected result

With $p = 0.5$: calculating an **upper bound how much** the MIA result **deviates from that of a random classifier**

Proposed Two-Stage Neural Network

GEFÖRDERT VOM Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union NextGenerationEU

German Research Center for Artificial Intelligence

Stage 1: Contrastive Learning with landmarks

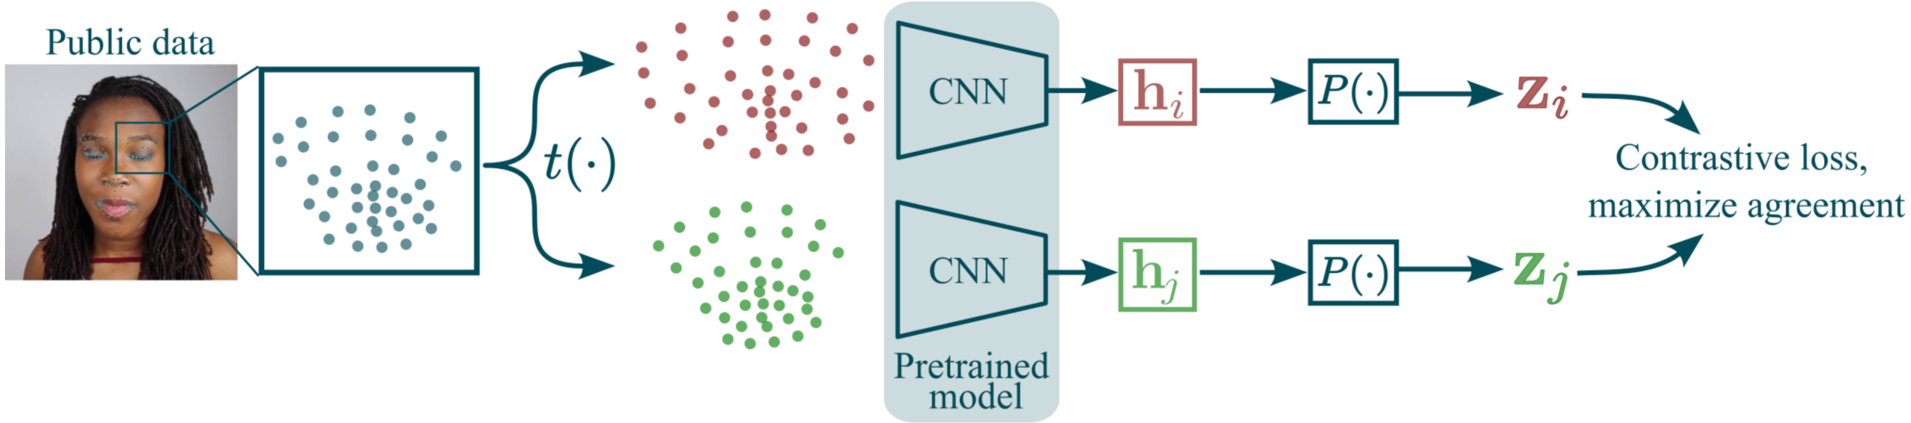

Photo from: P. Barros, N. Churamani, E. Lakomkin, H. Siqueira, A. Sutherland, and S. Wermter, "The OMG-Emotion Behavior Dataset," in 2018 International Joint Conference on Neural Networks (IJCNN), Pages: 7, Rio de Janeiro, Brazil: IEEE, Jul. 2018, pp. 1408–1414.

Proposed Two-Stage Neural Network

GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union
NextGenerationEU

German Research Center for Artificial Intelligence

Stage 1: Contrastive Learning with landmarks

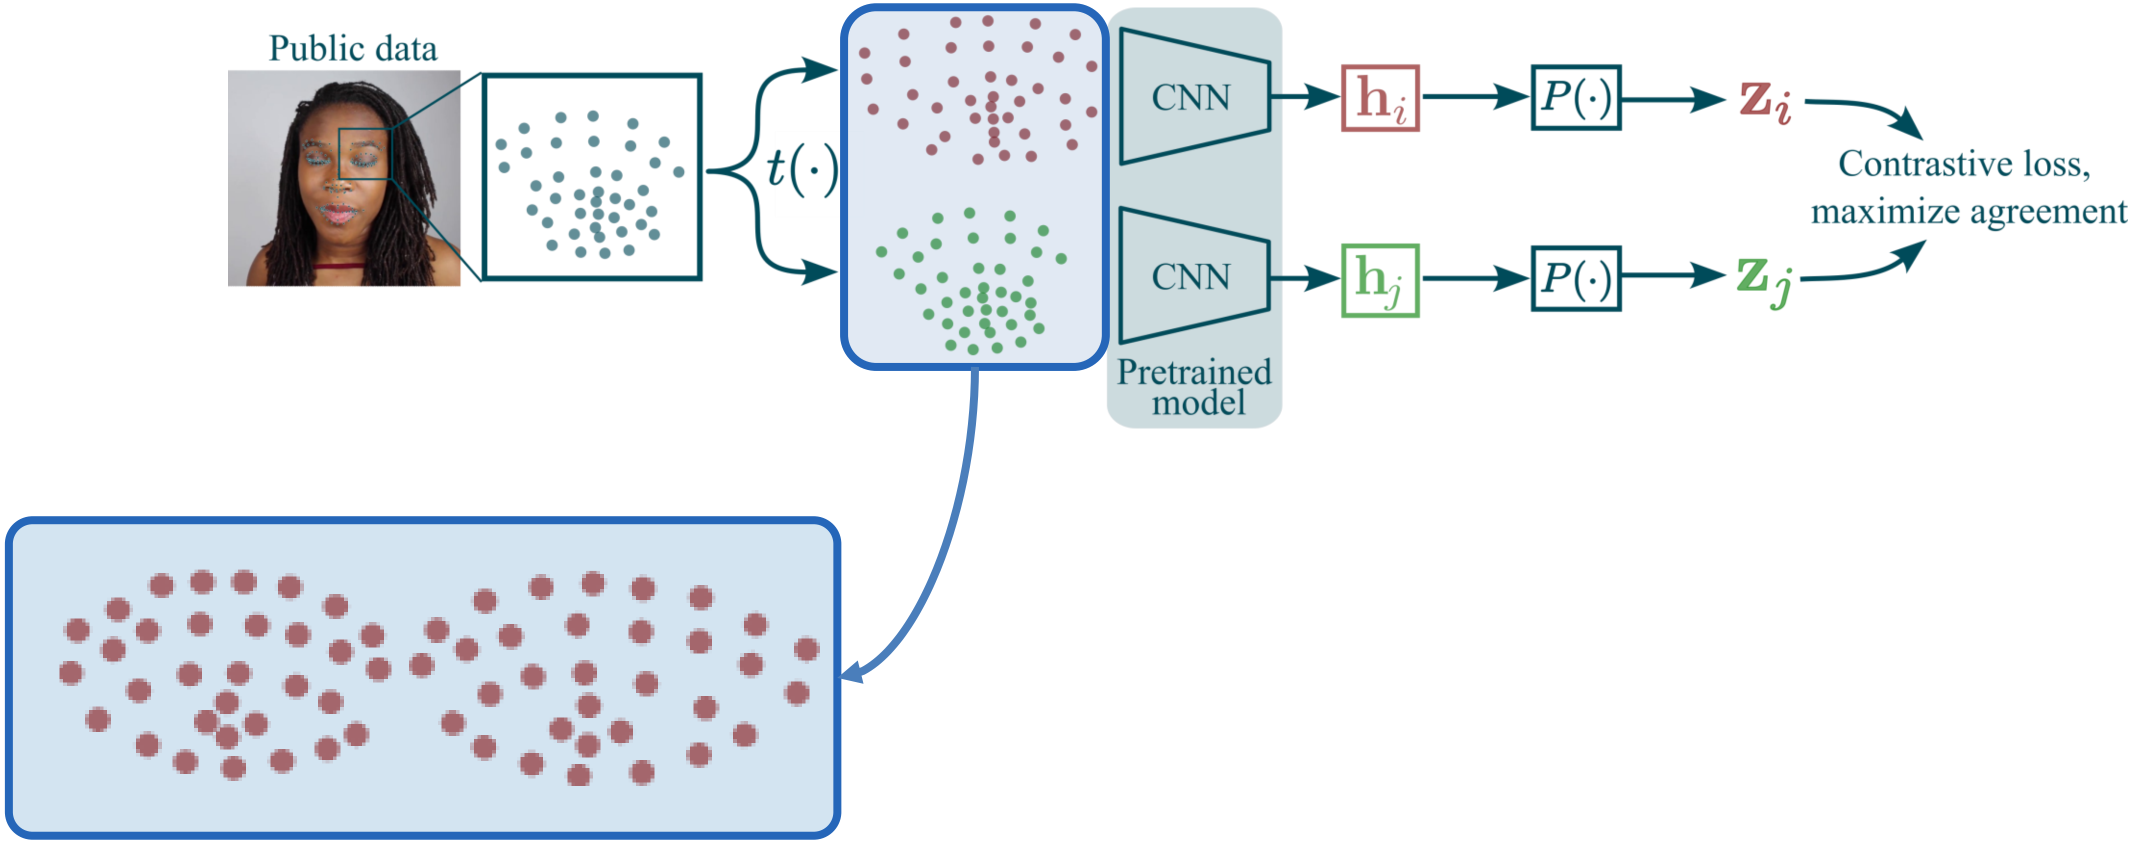

Proposed Two-Stage Neural Network

GEFÖRDERT VOM Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union NextGenerationEU

German Research Center for Artificial Intelligence

Stage 1: Contrastive Learning with landmarks

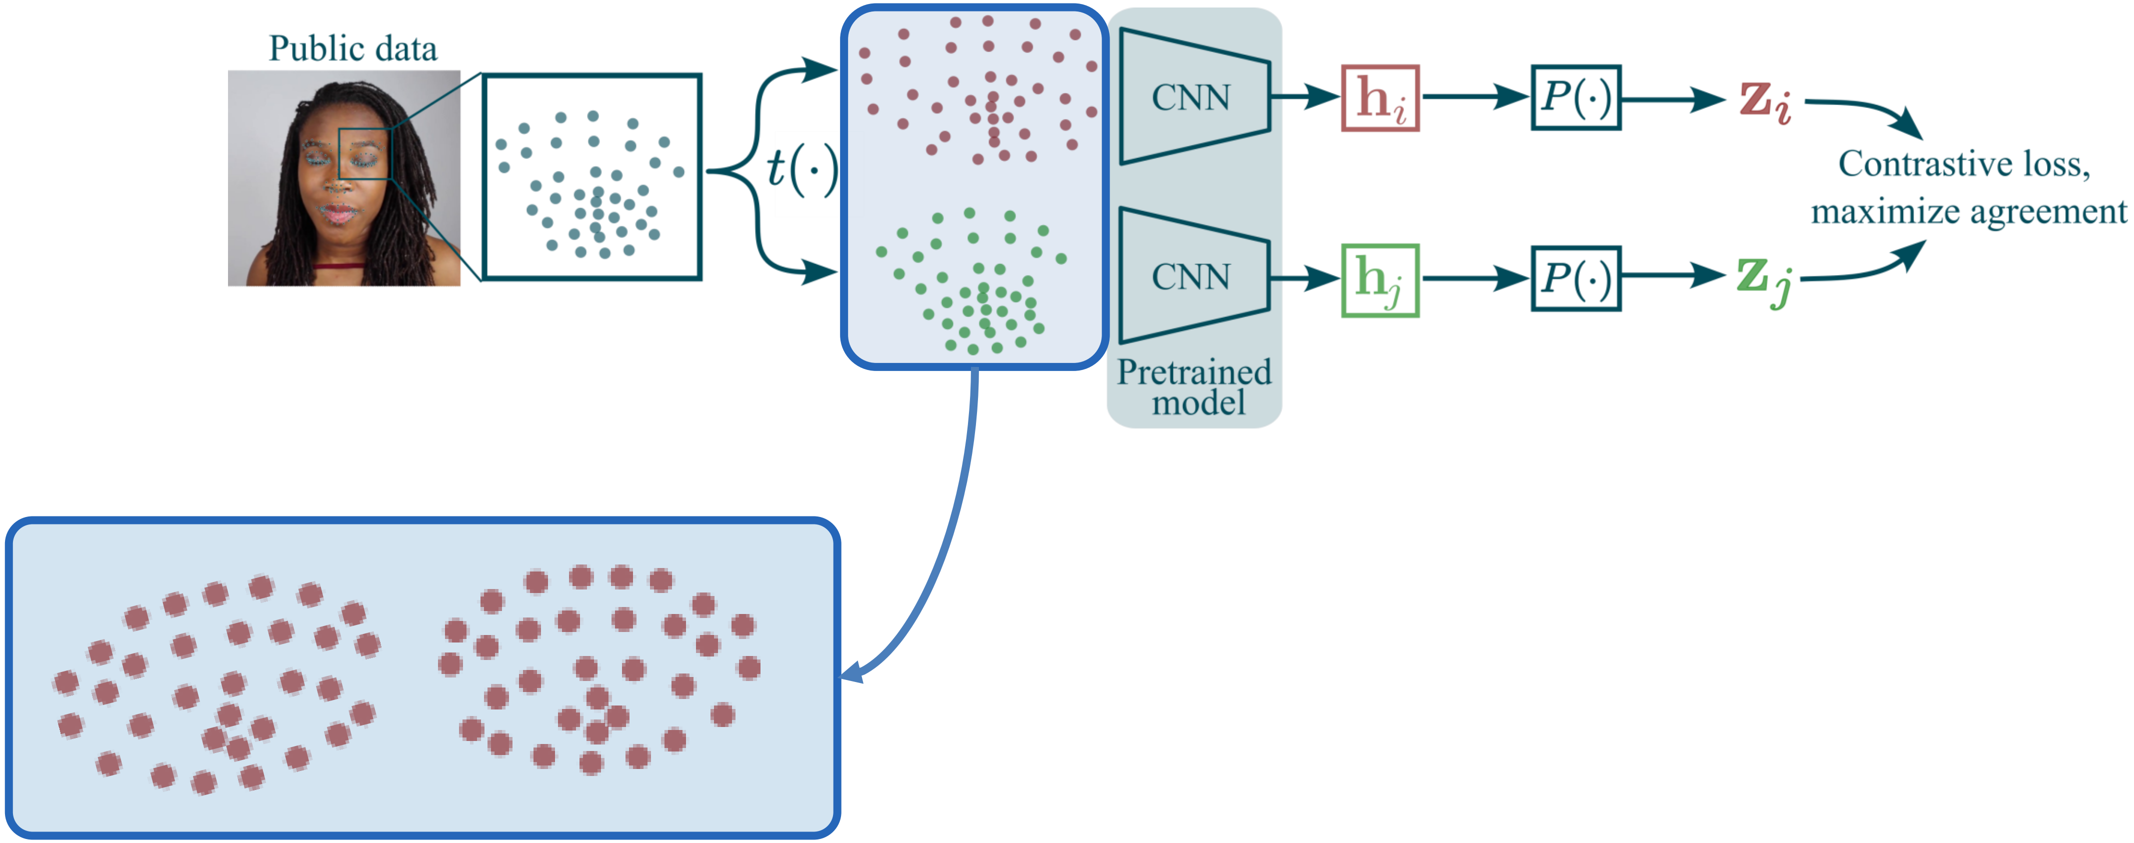

Proposed Two-Stage Neural Network

GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union
NextGenerationEU

German Research Center for Artificial Intelligence

Stage 1: Contrastive Learning with landmarks

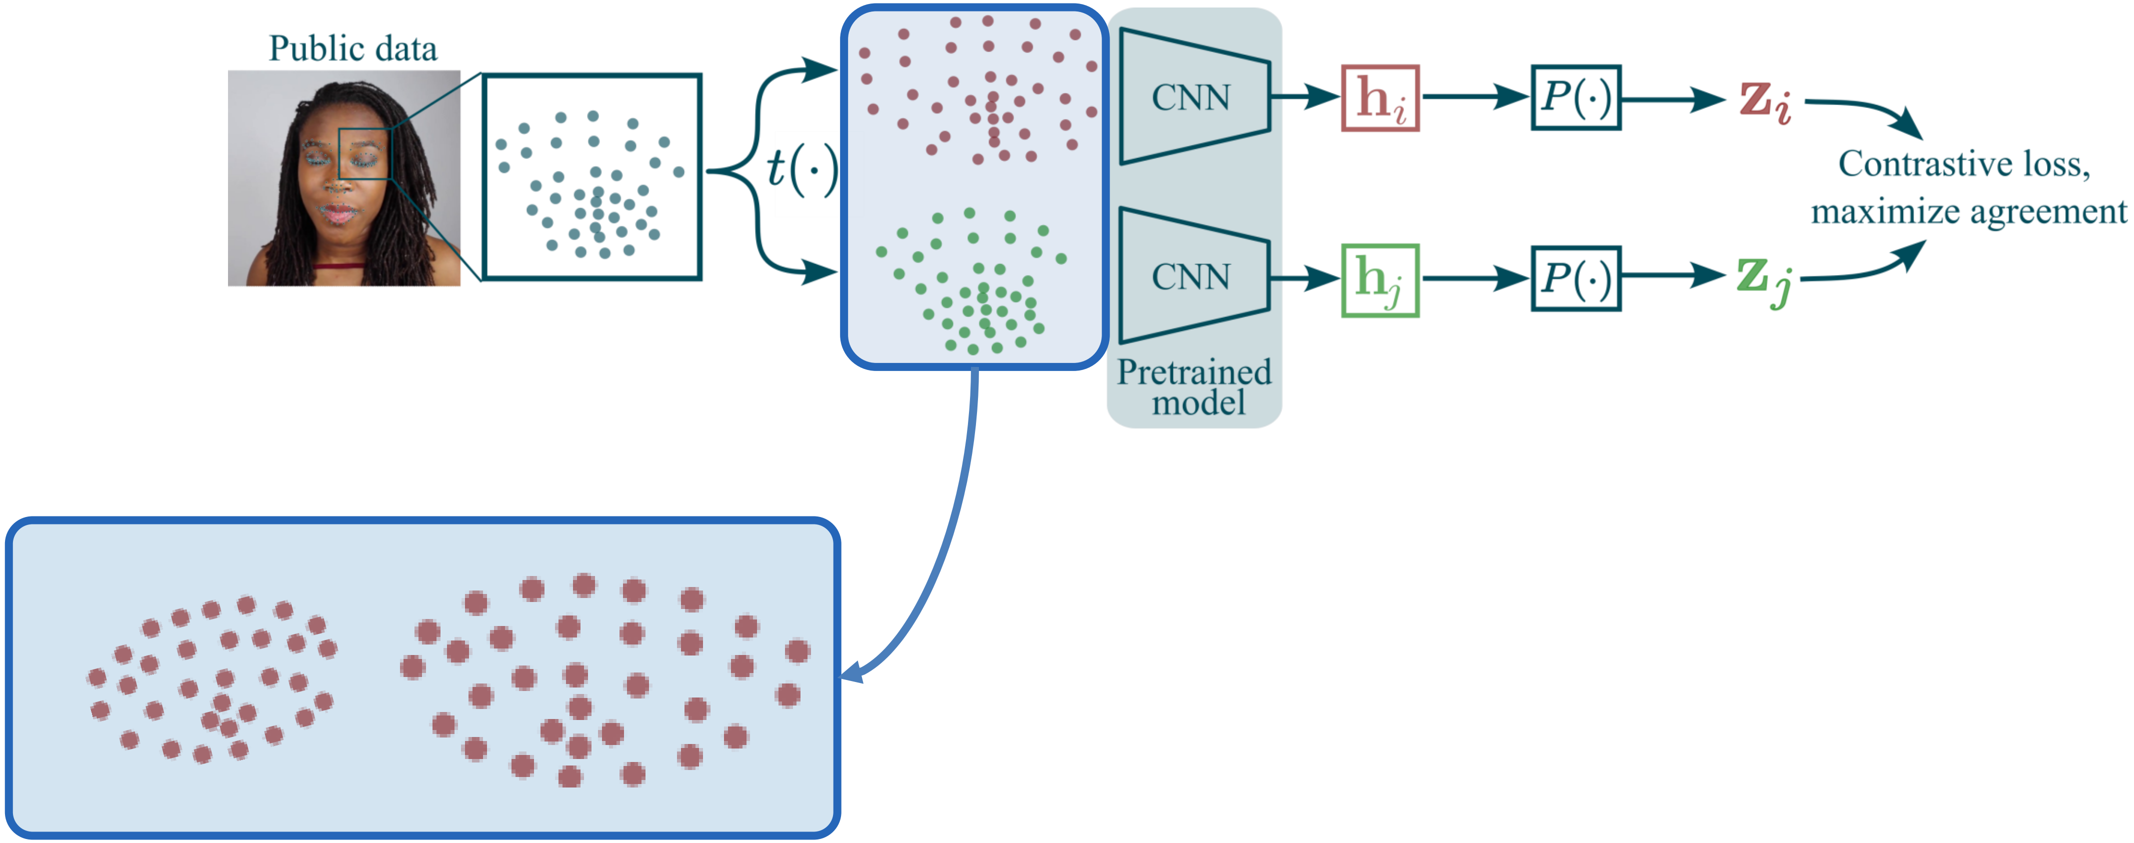

Proposed Two-Stage Neural Network

GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union
NextGenerationEU

German Research Center for Artificial Intelligence

Stage 1: Contrastive Learning with landmarks

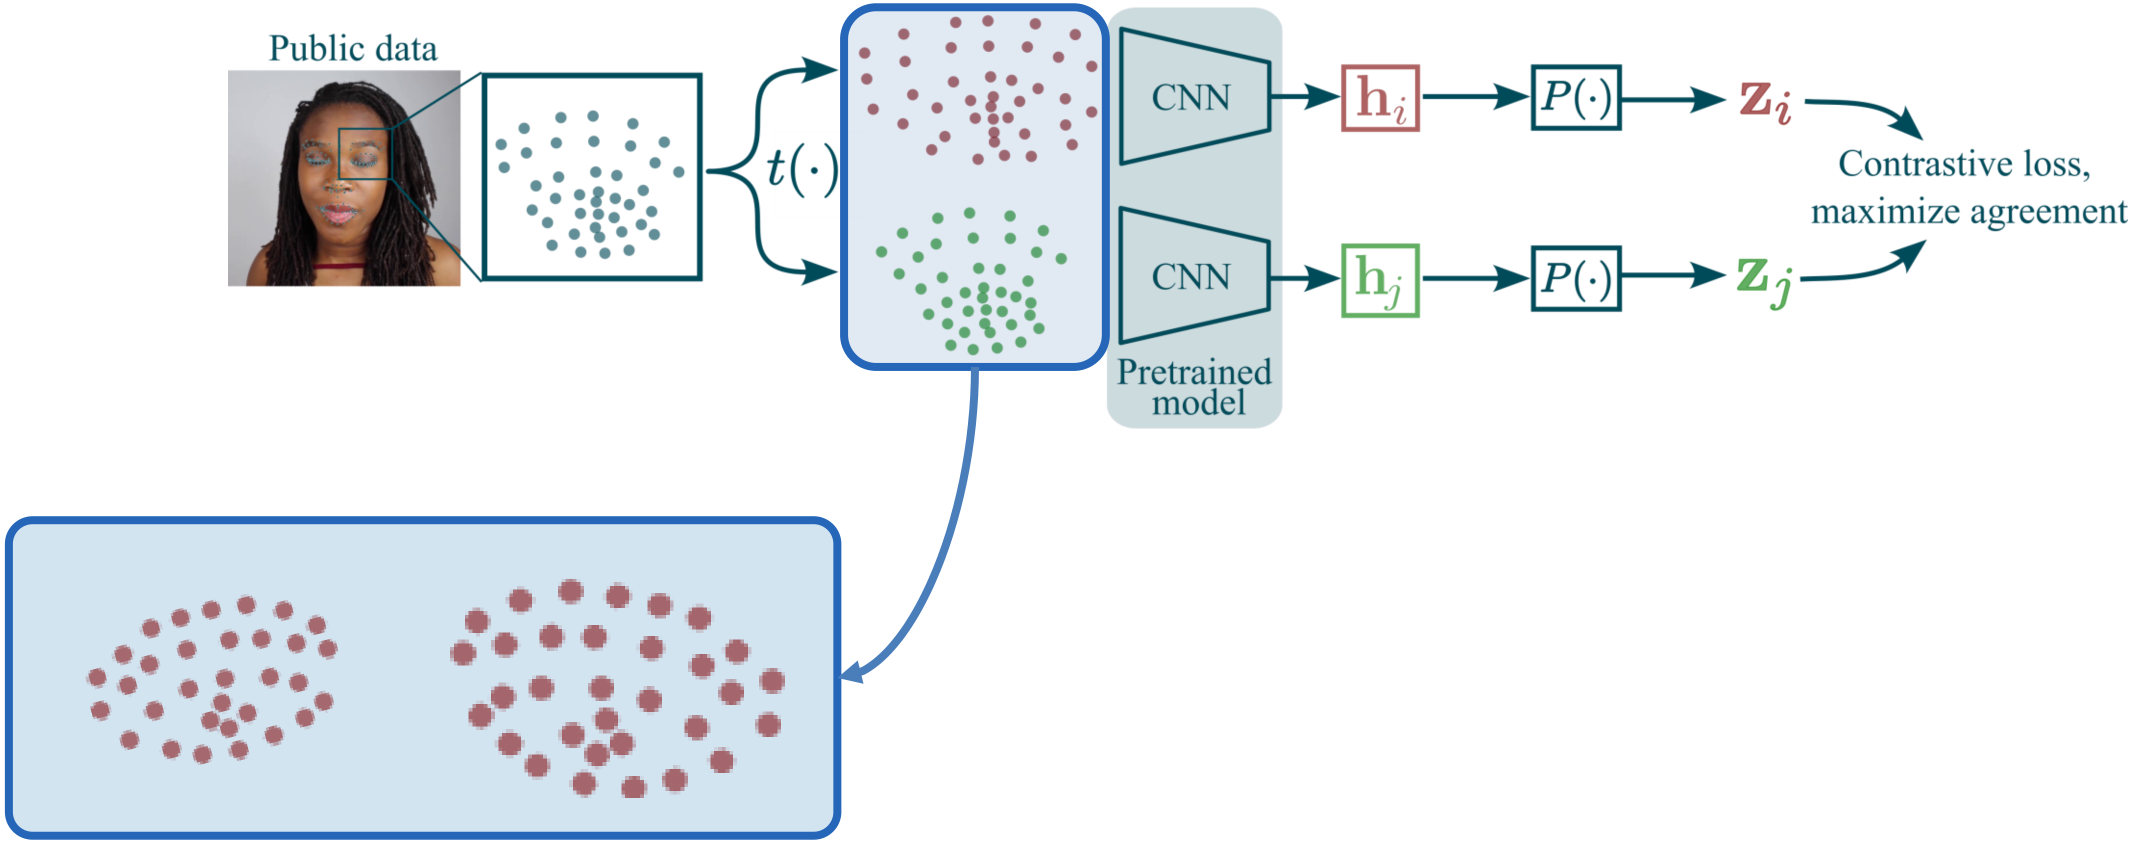

Proposed Two-Stage Neural Network

GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union
NextGenerationEU

German Research Center for Artificial Intelligence

Stage 1: Contrastive Learning with landmarks

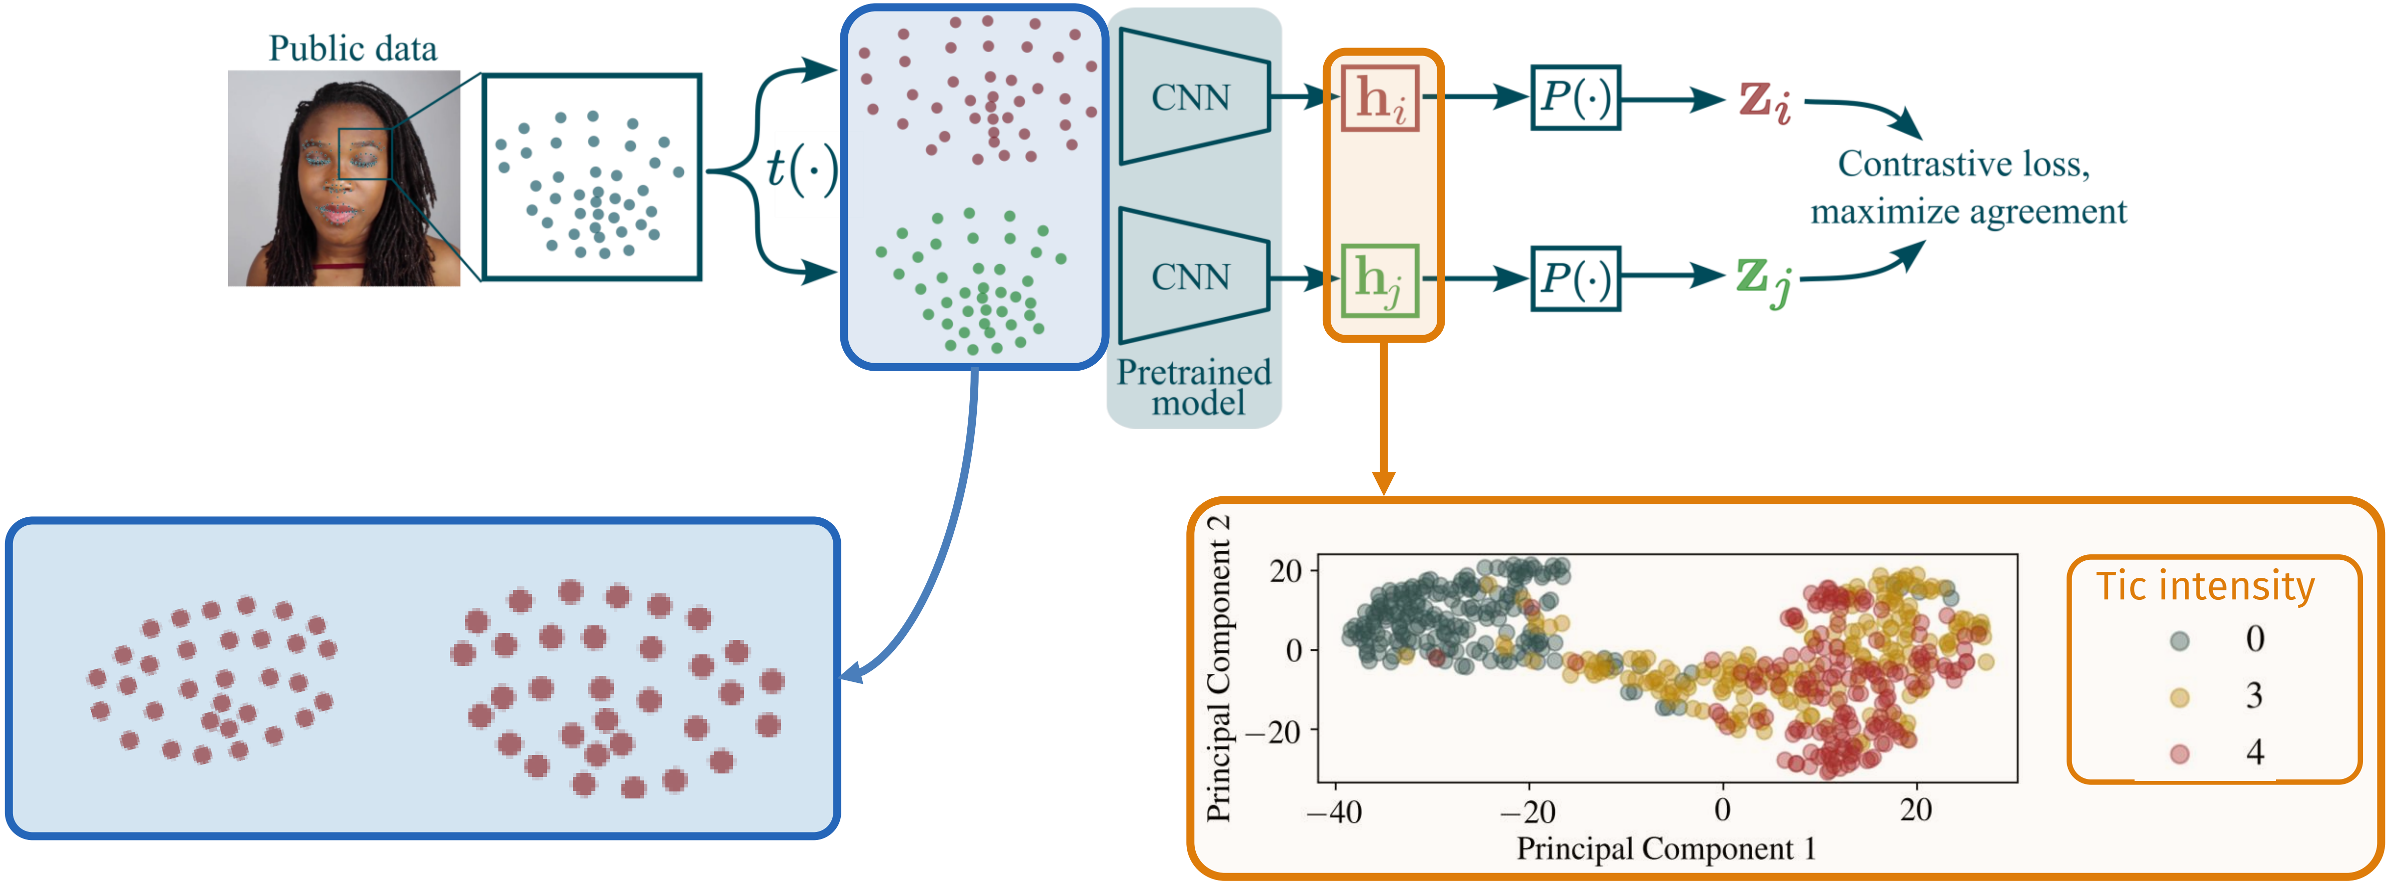

Proposed Two-Stage Neural Network

GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung
Finanziert von der Europäischen Union
NextGenerationEU
German Research Center for Artificial Intelligence

Stage 2: Training a classifier based on Contrastive Learning features

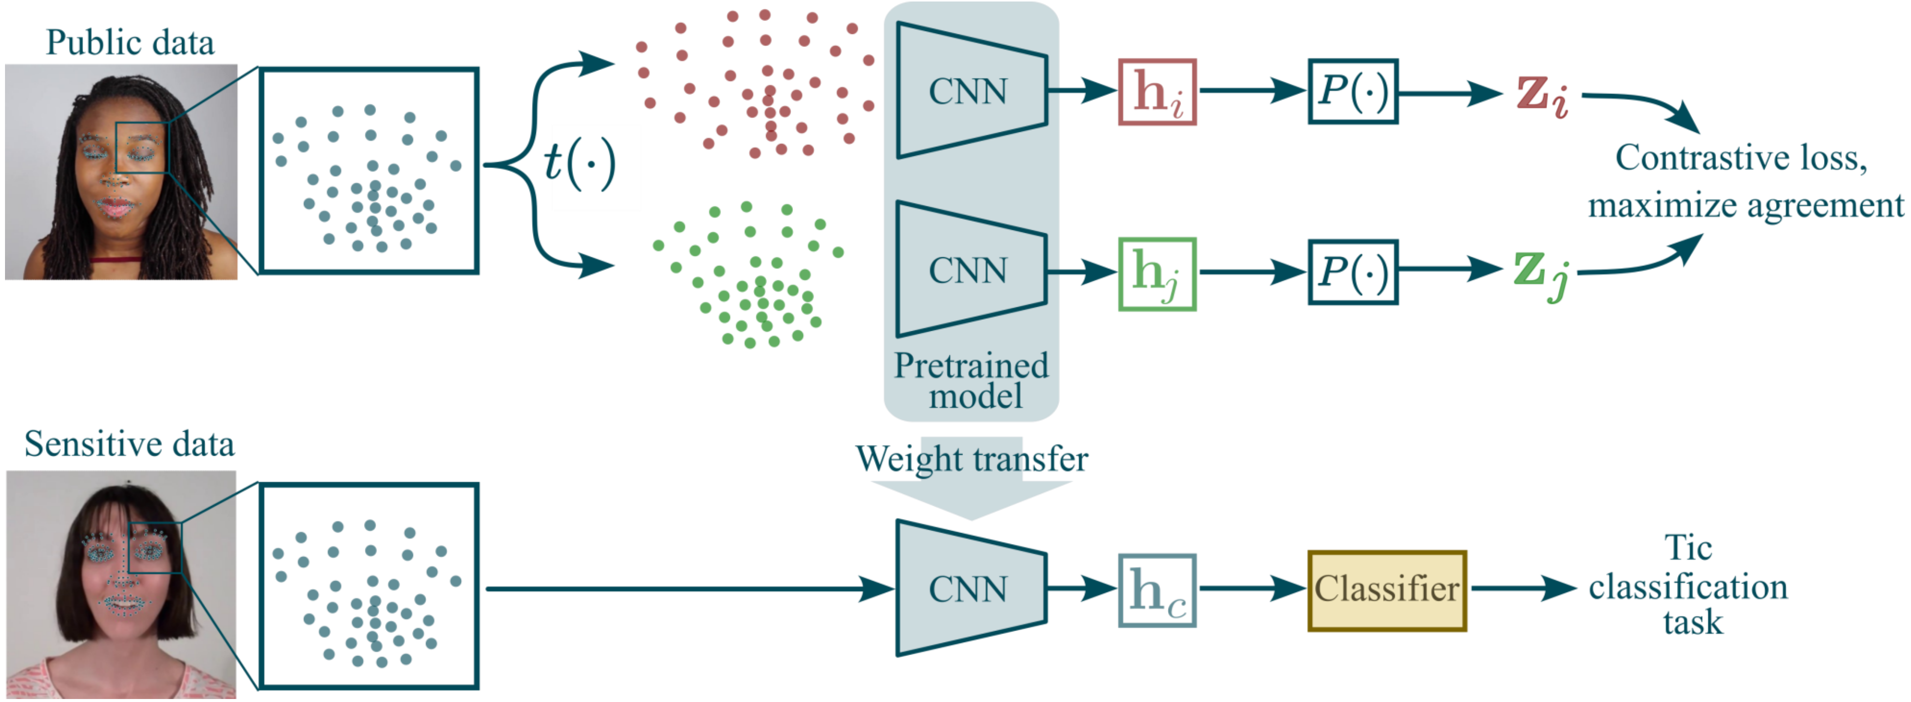

Photos from: P. Barros, N. Churamani, E. Lakomkin, H. Siqueira, A. Sutherland, and S. Wermter, "The OMG-Emotion Behavior Dataset," in 2018 International Joint Conference on Neural Networks (IJCNN), Pages: 7, Rio de Janeiro, Brazil: IEEE, Jul. 2018, pp. 1408–1414.

Tic Detection and Attack Results

GEFÖRDERT VOM Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union NextGenerationEU

German Research Center for Artificial Intelligence

Method	MIA Accuracy	Amplified MIA P[$random$]	Tic Detection Accuracy
Fully-Supervised CNN	73.10	4.07	81.92
Fully-supervised CNN + augment	65.65	23.00	85.76
Pretraining + Classifier	52.15	97.26	86.53
DP-Training of Classifier (ε=1)	50.14	99.80	80.09

(Amplified) membership inference attack (MIA) and tic detection accuracy for three different deep learning approaches.

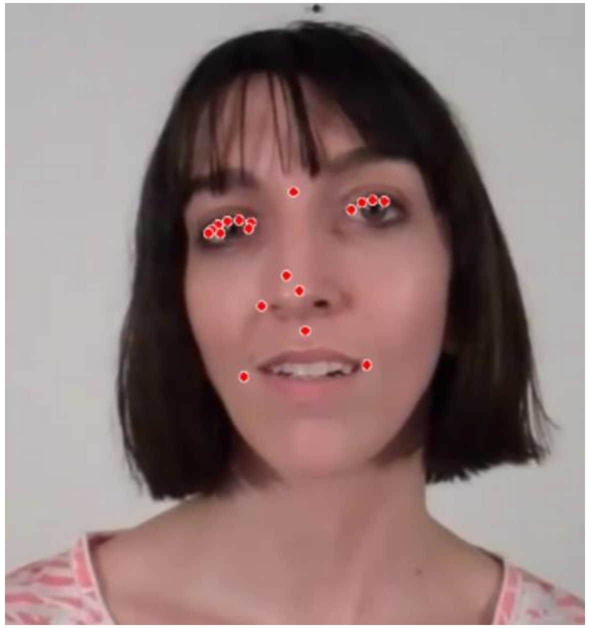

Most important landmarks determined via SHAP analyis

Conclusion

GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union
NextGenerationEU

German Research Center for Artificial Intelligence

- Supervised training on **small medical datasets**
  → **models leak information** about training data
- **Amplification of membership attacks** if multiple inputs per subjects are used
- **Two-stage learning approach leveraging data of healthy subjects**
  - Contrastive learning on landmarks
  - Effectively defends against MIAs
  - Reaches the highest classification accuracy of 86.5 %
  - Useful for various medical-related problems

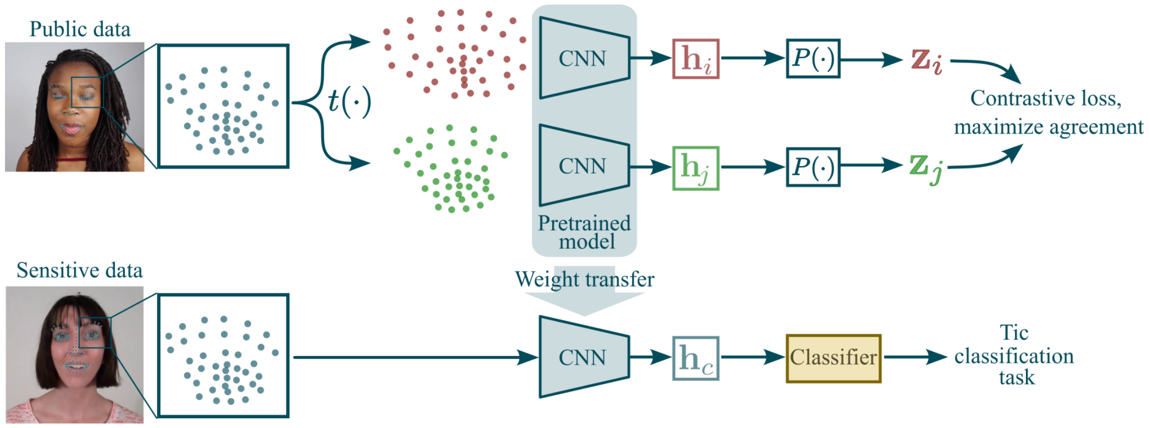

Past, Ongoing and Future Work

GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union
NextGenerationEU

German Research Center for Artificial Intelligence

Next steps:

- Provide provable privacy guarantees + and good utility also for small datasets
- Prepare AnoMed privacy challenge for video data

Semi-Supervised Training

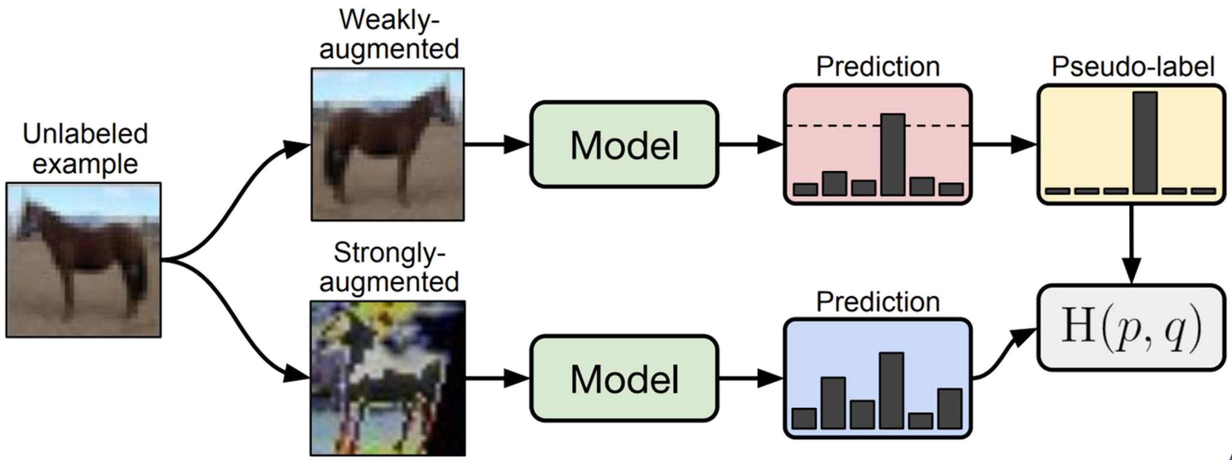

Differentially-private finetuning of large deep learning models

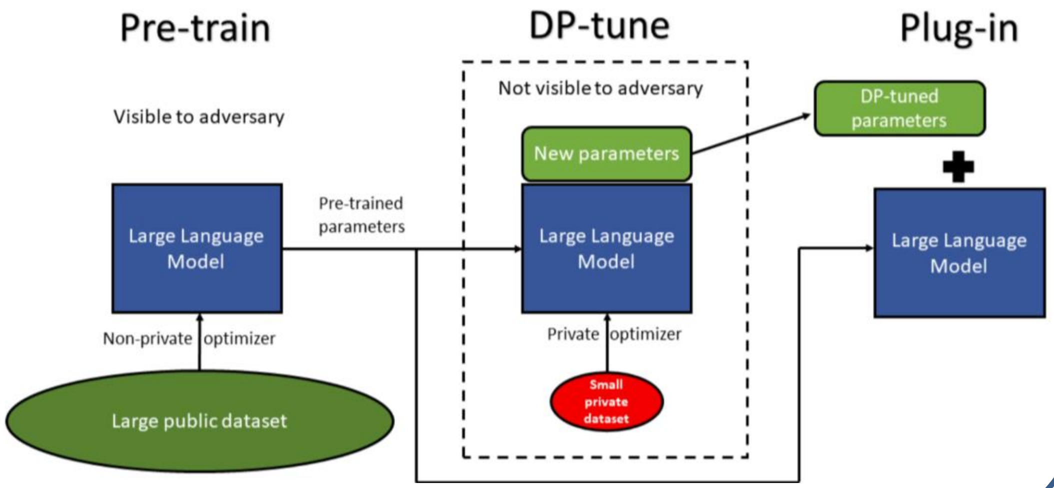

Past, Ongoing and Future Work

GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union
NextGenerationEU

German Research Center for Artificial Intelligence

Next steps:

- Provide provable privacy guarantees + and good utility also for small datasets
- Prepare AnoMed privacy challenge for video data

Semi-Supervised Training

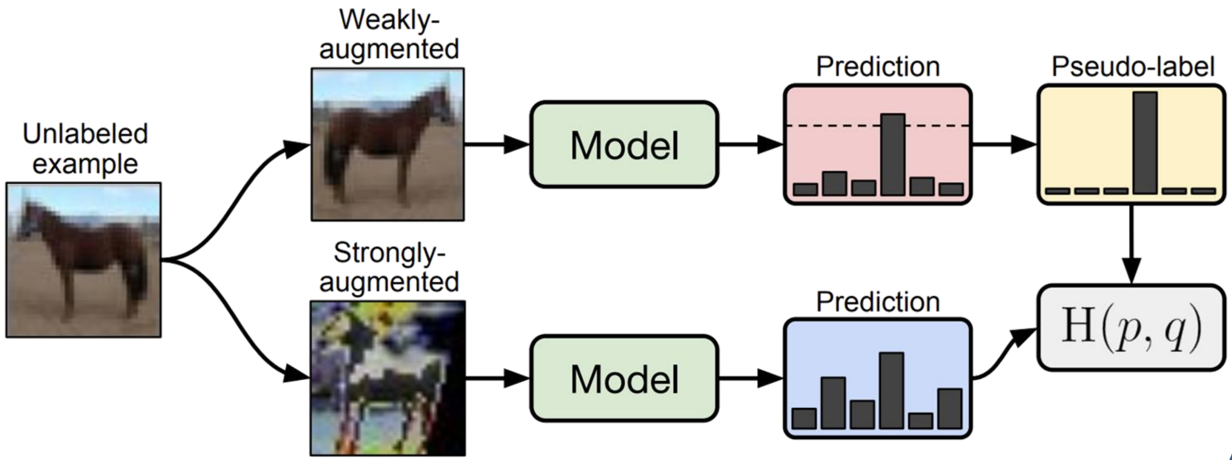

Differentially-private finetuning of large deep learning models

- Finetuning only the parameters that change the most during a gradient update
- Find the most efficient network update to save privacy budget
- Find collisions of two data points of different classes the in the neural network and finetune only this region
- **Finetune a low-rank approximation of the original network weights**

Past, Ongoing and Future Work

GEFÖRDERT VOM Bundesministerium für Bildung und Forschung

Finanziert von der Europäischen Union NextGenerationEU

German Research Center for Artificial Intelligence

Next steps:

- Provide provable privacy guarantees + and good utility also for small datasets
- Prepare AnoMed privacy challenge for video data

Semi-Supervised Training

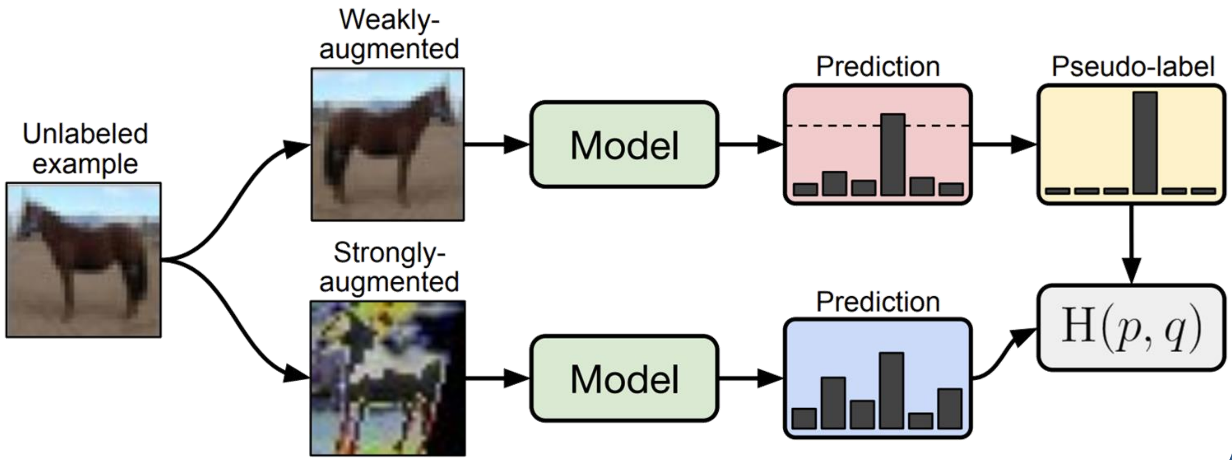

Differentially-private finetuning of large deep learning models

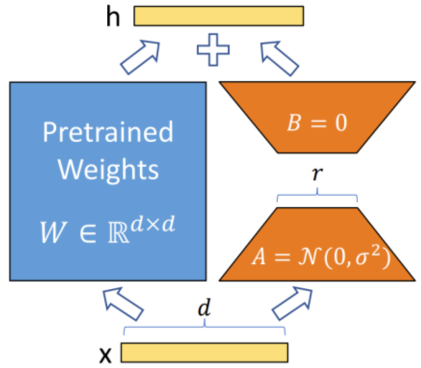

- **Finetune a low-rank approximation of the original network weights**

Contact

GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung
Finanziert von der Europäischen Union
NextGenerationEU
German Research Center for Artificial Intelligence

Nele Sophie Brügge

Researcher

DFKI Branch Office Lübeck

Artificial Intelligence in Medical Image Processing

Ratzeburger Allee 160
23562 Lübeck

nele.bruegge@dfki.de